arXiv:1511.06730v3 [stat.AP] 26 Sep 2016

A Bayesian hidden Markov mixture model to detect overexpressed chromosome regions.[1]

Vinícius Diniz Mayrink and Flávio Bambirra Gonçalves
Departamento de Estatística, Universidade Federal de Minas Gerais, Brazil.

Abstract

In this study, we propose a hidden Markov mixture model for the analysis of gene expression measurements mapped to chromosome locations. These expression values represent preprocessed light intensities observed in each probe of Affymetrix oligonucleotide arrays. Here, the algorithm BLAT is used to align thousands of probe sequences to each chromosome. The main goal is to identify genome regions associated with high expression values which define clusters composed by consecutive observations. The proposed model assumes a mixture distribution in which one of the components (the one with the highest expected value) is supposed to accommodate the overexpressed clusters. The model takes advantage of the serial structure of the data and uses the distance information between neighbours to infer about the existence of a Markov dependence. This dependence is crucially important in the detection of overexpressed regions. We propose and discuss a Markov chain Monte Carlo algorithm to fit the model. Finally, the proposed methodology is used to analyse five data sets representing three types of cancer (breast, ovarian and brain).

keywords: gene expression; microarray; Affymetrix; cancer; Gibbs sampling.

1 Introduction

DNA microarray platforms have been widely used in the past years to simultaneously measure the expression levels of a large number of genes; see Amaratunga et al. (2014) for a comprehensive coverage of recent advances in microarray data analysis. In particular, the high density Affymetrix GeneChip oligonucleotide array technology (Dalma-Weiszhausz et al., 2006) is the most popular platform used in biomedical researches focused on expression profiling and DNA analysis, at a genome global level or based on a subset of genes. Irizarry et al. (2003), Carvalho and Irizarry (2010), Mayrink and Lucas (2013), Li et al. (2013) and Mayrink and Lucas (2015) are few examples of works dealing with this type of data.

In the present study, we are concerned with the identification of chromosome regions associated with high expression measurements obtained from the aforementioned Affymetrix microarrays. In order to link each expression value to a location in the chromosome, we consider the sequence alignment algorithm BLAT (Kent, 2002) designed to map millions of sequence reads against the human genome. The main advantage of this algorithm when compared to others is its considerably faster performance, which allows for a more regular update of the genome. We provide further details about BLAT in Appendix A.

The identification of differentially expressed probe sets has been the topic of some recent works. In particular, Warren et al. (2007) proposed a method to generate detection calls "present", "marginal" and "absent" indicating the activity status of genes in the samples. In brief, the authors identify Affymetrix probe sets which cannot hybridize to the intended transcript, and use the empirical cumulative distribution of their intensities to derive a cutoff intensity. A probe set is thus classified as "Present" if its expression value is higher than the established threshold. The main critique to this method is the arbitrary choice of the cutoff point, which directly affects the classification of probe sets with intensities close to the unknown border separating the high and low categories.

In the present work, we are interested in the identification of overexpressed chromosome regions, based on the BLAT probe expression mapping. In particular, we analyse data from three types of cancer – breast,

[1]Accepted for publication in the Journal of the Royal Statistical Society - Series C (Applied Statistics).
See http://onlinelibrary.wiley.com/doi/10.1111/rssc.12178/full

Address for correspondence: Vinícius Mayrink, Departamento de Estatística, ICEx, UFMG. Av. Antônio Carlos, 6627, Belo Horizonte, MG, Brazil, 31270-901. E-mail: vdm@est.ufmg.br

ovarian and brain. We consider three data sets from the first type and one data set from each of the other two types. An overexpressed status suggests that a gene plays an important role in the progression of the tumor being investigated; the region identification provided in our study calls the attention for specific parts of the genome where one can find several genes (not only those from the arrays) potentially having a key contribution for understanding the disease.

In order to avoid choosing an arbitrary threshold to classify a value as overexpressed, we take advantage of the irregularly spaced serial structure of the data by using the neighbourhood and distance information. More specifically, we propose a semiparametric hidden Markov model to fit the observations using a mixture distribution with some gammas and a single Gaussian components. The latter is supposed to accommodate the overexpressed clusters. The gamma components are inserted to deal with the skewness and multimodality exhibited by the data. Finally, we consider a mixture of discrete distributions to model the uncertainty on the Markov dependence.

We develop a Markov chain Monte Carlo (MCMC) algorithm to perform inference under the Bayesian paradigm. The algorithm is a Gibbs sampling designed to have good convergence properties. This is achieved by a careful choice of the blocking and update schemes. The most crucial step of the algorithm involves a backward-filtering-forward-sampling strategy.

Mixture models for microarray data have been explored in several previous works. For example, Lewin et al. (2007) present a Bayesian hierarchical model to determine the expression status of genes using a mixture prior distribution for the parameters representing differential effects. Broet et al. (2002) consider a mixture of Gaussians at the data level of the model hierarchy to identify expression changes. A Gaussian mixture is also used in Do et al. (2005) for the analysis of differentially expressed genes between normal and colon cancer tissue samples. Some other examples are: Efron et al. (2001), Newton et al. (2001), Parmigiani et al. (2002), Broet et al. (2004) and Dean and Raftery (2005). None of these works, however, consider the information regarding the position of the expressions along the chromosomes.

In computational biology, one can easily find studies evaluating data sets containing genome map position information (see Yi et al., 2005; Xu et al., 2013; Cheung et al., 2005; Lucas et al., 2010). In the context of copy number data, Pollack et al. (2002) explores genome-wide measurements of DNA copy number alteration by array CGH; their data are $\log_2$-base fluorescence ratios images depicting amplifications, deletions and unchanged values across the chromosomes. Broet and Richardson (2006) propose a novel method called CGHmix to investigate copy number changes, based on a spatially structured mixture model. Other few examples accounting for the spatial dependency along the chromosomes are: Autio et al. (2003), Jong et al. (2004) and Picard et al. (2005). Another common research topic using chromosome location data is the study or comparison of genome alignment tools; for example, Allen et al. (2011) review and compare different approaches to map probes across different microarray platforms.

There is an extensive literature dedicated to methodologies to deal with irregularly spaced serial data. In a continuous-time context, discretely observed diffusion processes might be one option (see Beskos et al., 2006). Many studies, including the present paper, assume a discrete-time configuration which can be appropriately explored through a particular case of a more general geostatistical model with neighbourhood weights decaying with distance; see, for example, dynamic and state-space models investigated in Shephard (1994), Fruhwirth-Schnatter (1994) and Gelfand et al. (2005).

One-dimensional change point models may be considered to address the clustering problem accounting for the serial structure of the data. These methods were developed to identify parameter shifts in a distribution assumed for the sequential data. An interesting approach is the Product Partition Model (PPM) introduced by Hartigan (1990) and Barry and Hartigan (1992); its spatial version, presented in Page and Quintana (2015), uses the Euclidean distances between neighbours to determine the probability of the partitions. However, the model fit would be burdensome to handle thousands of chromosome locations; the PPM generates the unknown change points and this may not be a practical task. As an alternative, Chib (1998) proposed a model assuming a fixed number of regimes and treating the breaking process as a Markov chain with transition probabilities constrained so that regimes come in a non-reversible sequence. The author indicates that the transition probabilities can be a function of covariates (e.g. the distance information); however, the Markov dependence is assumed true across the whole series despite the distance between neighbours.

The outline of this paper is as follows: in Section 2, we present all aspects considered to obtain the final data set for analysis; they include: short description of the microarray structure and a scale transformation

to better manage the distances between observations. In Section 3, we propose a hidden Markov mixture model to identify overexpressed regions of the chromosomes. Section 4 presents and details the Bayesian inference procedure, including an MCMC algorithm and an strategy to perform the cluster detection based on the output of this algorithm. Section 5 presents the analysis of five data sets representing three types of cancer. Finally, in Section 6, we summarise and discuss the main conclusions of the study.

2 The data

We consider Affymetrix HG-U133A oligonucleotide arrays previously explored in Miller et al. (2005), Wang et al. (2005), Sotiriou et al. (2006), Marks et al. (1991) and Freije et al. (2004). The first three studies are related to breast cancer; the last two refer to ovarian and brain cancer, respectively. Hereafter, we will denote these data sets by "Breast 1" (251 samples), "Breast 2" (286 samples), "Breast 3" (189 samples), "Ovarian" (141 samples) and "Brain" (59 samples); these arrays contain expressions of thousands of genes. In brief, an Affymetrix GeneChip consists of a quartz wafer (chip) to which are attached approximately 500,000 different known 25-mer oligonucleotides. The data is obtained through an hybridization procedure where: mRNA are extracted from the cells or tissues of interest, labeled with fluorescent tags and combined with the chip. The mRNA single strands are expected to connect to its complementary sequences, if found, in a specific spot within the array. These array spots are called "probes" and they represent a fraction of a gene. Each transcript is represented on an array by a series of 11-20 probe pairs known as a probe set; for simplicity, we may refer to a probe set as a gene. Each pair consists of a perfect match probe (PM), with its 25-base sequence identical to the gene of interest, and a mismatch probe (MM), whose sequence is the same as the PM except for the 13^{th} position, where the base is set to the PM complementary. The MM probe was introduced by Affymetrix as a measure of non-specific binding or cross-hybridization. After the hybridization, the array is washed to remove unconnected material and a laser is applied to activate the fluorescence. Finally, a scanner is used to measure a positive and continuous value of light intensity representing the expression signature in each probe.

In the proposed application, these expression data are mapped to chromosome locations through the alignment tool BLAT – see Appendix A. Another step to be considered for the raw Affymetrix expression data is a preprocessing routine to remove part of the systematic noise effect aggregated during the experiment to build the arrays. We consider the most popular method known as Robust Multi-chip Average (RMA) (see Irizarry et al., 2003), also described in Appendix A.

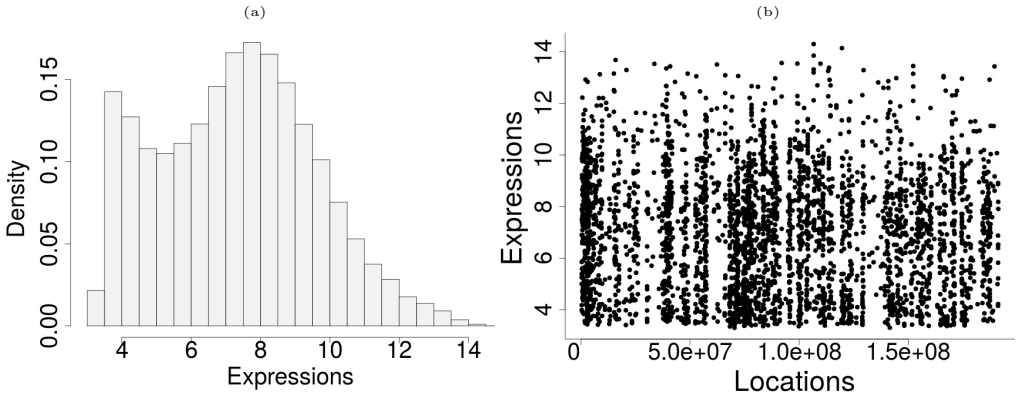

Figure 1: Graphs related to the data set "Breast 1". Panel (a): Histogram displaying the distribution of the medians of the preprocessed $\log_2$-base intensities identified via BLAT in all chromosomes (1-22, X, Y). Panel (b): points showing the positions of the medians along chromosome 4.

Figure 1 presents a histogram in panel (a) indicating a skewed and multimodal distribution for the expression data (Breast 1) mapped to the chromosomes. Panel (b) shows the spatial configuration of the data along the chromosome 4 with distances varying between data points. Here, we consider the original scale of chromosome positions identified via BLAT; the minimum, average and maximum distances are: 1, 53,300 and 3,160,573, respectively. The total number of probes mapped to chromosome 4 is 3,582. Other

chromosomes in the human genome have similar results, but their lengths differ; in particular, Y is the shortest one.

The analysis of Figure 1 (b) motivates the use of a hypothesis test to verify the spatial dependence in the data. A well known test for this purpose is based on the Moran's I statistics (Moran, 1950). We consider the `R` package `spdep` (Bivand and Piras, 2015) to apply a Monte Carlo version of this test. In brief, the Moran's I, with weights based on the inverse distance between locations, is calculated for the observed data and for many random permutations of the expression values. The resulting p-value represents how often the observed Moran's I is close to those obtained for the permuted data with no spatial association.

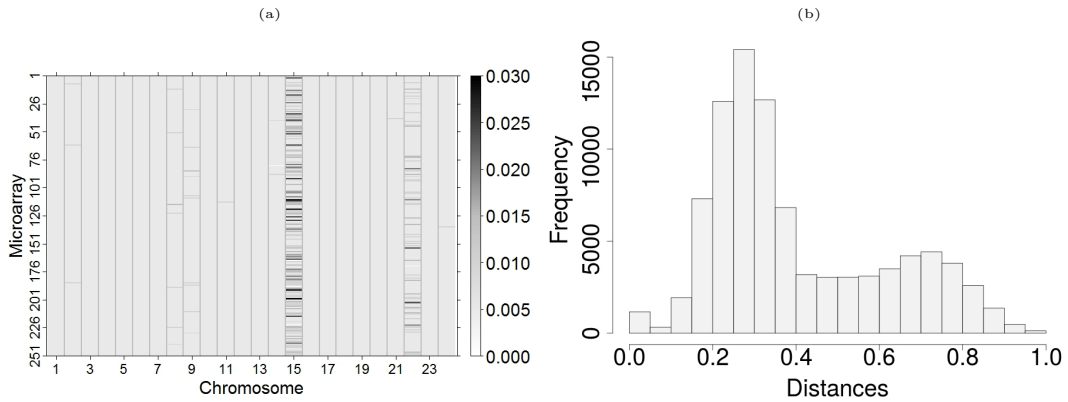

Figure 2: Panel (a): Image displaying all p-values of the Moran's I permutation test applied to the data set "Breast 1"; each row represents a microarray and each column represents a chromosome (1-22, X and Y). Panel (b): Histogram of all rescaled distances between locations in all chromosomes.

Figure 2 (a) shows a heat map graph representing the magnitude of the Moran's I permutation test p-values for each chromosome and each microarray of the "Breast 1" data; similar results were obtained for the other data sets. Note that all p-values are small (< 0.05) suggesting that the spatial association is significant. In particular, chromosome 15 concentrates the largest p-values, but none of them are high enough to reject the association. In summary, these tests indicate presence of spatial dependence in the aligned data and thus motivate a model taking into account the distances between neighbours.

Note that the original BLAT distances are too large and might lead to computational difficulties to fit a model. In order to avoid this issue, we rescale those distances to the interval (0,1) using two steps: calculate the log-distances and divide each of them by the maximum log-distance. The model proposed in the next section will assume the distances as covariates, therefore, other transformations can be considered and this might affect the inference, but it does not change the structure of our model. Taking the logarithm will preserve high differences among the smallest values and practically even out the highest values. This is coherent with the model we propose to perform the analysis, as it will be discussed in the next section. Figure 2 (b) shows a histogram for all rescaled distances in the chromosomes.

3 Hidden Markov mixture model

Let n be the number of locations to be analysed and L the number of arrays. Define $X = \{X_i,\ i = 1, \ldots, n\}$, $Y = \{Y_i,\ i = 1, \ldots, n\}$ and $Y_i = \{Y_{il},\ l = 1, \ldots, L\}$, where X_i is the true expression (latent) of location i and Y_{il} is the preprocessed expression of location i observed with noise (w.r.t. X_i) in array l.

Since the distribution of each Y_i is unimodal, fairly symmetric and L is large (59-286 in our applications), we assume $X_i = \tilde{Y}_i$, for simplicity, where $\tilde{Y}_i$ is the median of Y_i. Nevertheless, the variability of the replications could be modeled by assuming, for example, $Y_{il} \sim N(X_i, \tau_i^2)$.

The main goal of this analysis is to detect overexpressed regions in each chromosome. The naive attempt of fixing an arbitrary threshold value and looking at the probes with expression intensity higher than this value is discarded for several reasons. Firstly, there is no explicit gap in the distribution of the expression intensity values in each chromosome. Secondly, there is no clear scale of the measured intensity, due to all the preprocessing of the data. Finally, this attempt would not take into account the dependence known to

exist among near probes. It is natural to assume that near probes, in the chromosome sequence, are likely to have similar expression intensities. As an example, see the identification of chromosomal regions with DNA copy number alteration in Pollack et al. (2002) and Lucas et al. (2010).

Our approach considers an stochastic model that uses the dependence among near probes to detect the overexpressed regions. More specifically, the regions of interest are linked to the component of highest mean in a mixture distribution. This way, no fixed threshold needs to be provided. The dependency information is used in a way that near probes are likely to be in the same component of the mixture. Generally speaking, the overexpressed regions are detected as sequences of probes that have the highest expression intensities in the chromosome. The dependence structure disfavors the detection of spikes, i.e. an isolated probe and none of its neighbours being detected as overexpressed. Finally, informative prior distributions for the parameters of the proposed model are crucial to achieve the desired goal.

The multimodality in Figure 1 (a) motivates the use of a mixture distribution to model the expressions X. This approach also allows the definition of the overexpressed regions, in a stochastic way, by associating them to one of the components in the mixture. This is the component with the highest mean and it is assumed to be a normal distribution. The remaining components of the mixture are assumed to be gamma distributions, which is motivated by the positive, multimodal and skewed features of the data in Figure 1 (a) (see Wiper et al., 2001). One might consider here a mixture of normal distributions instead; however, this would require a larger number of components to fit the data (see Baudry et al., 2010). Mixtures of gammas have been explored in different contexts (see, for example, Dey et al., 1995; Nascimento et al., 2012).

Given that BLAT maps thousands of probe expressions to each chromosome and that the majority of locations have neighbours relatively close, it is unlikely that a single location with an intensity much higher than its neighbours is in fact an observation within an overexpressed region. In other words, the neighbourhood structure is crucial in the detection of such regions. This spiked configuration suggests a local atypical expression and is expected to be modelled (associated) to the right tail of a gamma component in the mixture. The local outlier may be explained by cross-hybridization, incorrect probe mapping and other aspects involving the microarray assembling and the alignment tool. The cases where an isolated high expression probe may belong to a potential overexpressed region are those where its nearest observed neighbours are too far away.

We use a Markov structure to model the dependence among probes. However, given the irregular structure of distances between locations in the data, a discrete time Markov chain on (a discretised version of) the distance scale would be both algebraically and computationally expensive. Therefore, we define a Markov structure on the ordering of the probes, i.e. the Markov dependence (if existent) is the same for any consecutive pair of probes. However, such dependence is assumed to be uncertain, i.e. to exist with some probability. In a lower level of the model, we use the distance information to model this probability.

Let K be the number of gamma components in the mixture. This means that the full mixture has $K+1$ components where the one with the largest mean is normal. Define $q_0 = (q_{01}, \ldots, q_{0(K+1)})'$ as a probability vector and $Q = \{q_{k_1 k_2}\}$, $k_1, k_2 = 1, \ldots, K+1$, as a transition matrix of a $K+1$ states and discrete time Markov chain. Let also q_k, $k = 1, \ldots, K+1$, be the k-th row of Q. Let F_k be the c.d.f. of a gamma distribution with mean θ_k and shape parameter η_k, for $k = 1, \ldots, K$, and F_{K+1} is the c.d.f. of a normal distribution with mean μ and variance σ^2. Denote f_k as the density implied by F_k, $k = 1, \ldots, K+1$. Define the vector $Z_i = (Z_{i,1}, \ldots, Z_{i,K+1})'$ such that $Z_{i,k} = 1$ indicates that X_i belongs to the k-th mixture component, 0 otherwise. We propose the following model:

$$
\begin{aligned}
(X_i | Z_{i,k} = 1) &\sim F_k, \;\; i = 1, \ldots, n, \text{ all independent;} && (1)\\
(Z_1 | q_0) &\sim \text{Mult}(1, q_0); && (2)\\
(Z_i | Z_{i-1,k} = 1, \rho_i, q_0, Q) &\sim (1-\rho_i)\,\text{Mult}(1, q_0) + \rho_i\,\text{Mult}(1, q_k), \;\; i = 2, \ldots, n, && (3)
\end{aligned}
$$

where Mult refers to the multinomial distribution and the parameters indexing the mixture components are supressed from the conditional notation. The distribution in (3) states that a Markov dependence between the expressions in the locations $i-1$ and i is present with probability ρ_i.

The model in (1)-(3) is not an ordinary mixture model because the X_i variables are not marginally (w.r.t. Z) independent. Their dependence, however, is defined in a second (and latent) level of the model through a Markov structure. This fact qualifies our model as a Hidden Markov Model (HMM). For modelling and computational reasons, we define Bernoulli random variables W_i, $i = 2, \ldots, n$, that indicate the presence of

the Markov dependence. That is,

$$\begin{aligned}
(Z_i|Z_{i-1,k}=1, W_i=0, q_0) &\sim \text{Mult}(1, q_0), \;\; i=2,\ldots,n; &(4)\\
(Z_i|Z_{i-1,k}=1, W_i=1, Q) &\sim \text{Mult}(1, q_k), \;\; i=2,\ldots,n; &(5)\\
(W_i|\rho_i) &\sim \text{Ber}(\rho_i), \;\; i=2,\ldots,n. &(6)
\end{aligned}$$

For notation reasons, define $W_1 = 0$ almost surely and $(Z_1|Z_0, W_1 = 0, q_0) := (Z_1|q_0)$.

Our Bayesian model is fully specified by adopting appropriate prior distributions. Let Φ be the c.d.f. of the standard normal distribution and define d_i as the (transformed) distance between locations $i-1$ and i. We adopt the following prior specifications:

$$\begin{aligned}
q_0 &\sim \text{Dir}(r_0);\\
q_k &\sim \text{Dir}(r_k), \text{ all independent;}\\
\theta_k &\sim IG(t_{1k}, t_{2k}), \; \theta_1 < \ldots < \theta_K;\\
\eta_k &\sim G(e_{1k}, e_{2k}), \text{ all independent;} &(7)\\
(\mu, \sigma^2) &\sim NIG(m, v, s_1, s_2), \; \mu > \theta_K;\\
(\rho_i|\beta) &= \Phi(\beta_0 + \beta_1 d_i);\\
\beta = (\beta_0, \beta_1)' &\sim N_2(\mu_0, \Sigma_0).
\end{aligned}$$

In the specifications above, Dir, IG, G, NIG and N_2 refer to the Dirichlet, inverse gamma, gamma, normal-inverse-gamma and bivariate normal distributions, respectively. Assume the following parametrisation: shape t_{1k} and scale t_{2k} in the inverse gamma prior for θ_k, and mean e_{1k} and shape e_{2k} in the gamma prior for η_k. Additional notation: $Z = \{Z_i, \; i = 1,\ldots,n\}$, $W = \{W_i, \; i = 1,\ldots,n\}$, $\theta = \{\theta_k, \; k = 1,\ldots,K\}$ and $\eta = \{\eta_k, \; k = 1,\ldots,K\}$.

The Markov dependence of the probes is expressed in terms of the mixture component indicator variable Z_i. The existence of this dependence is stochastically explained by the distance between the probes through the probit regression for the W_i variables. Vector q_0 could be defined as the stationary distribution of the Markov chain with transition matrix Q. However, we do not make this restriction in the model as it would make the inference methodology much harder (more specifically, the sampling step of the q_k's).

The normal component of the mixture is expected to accommodate the expressions of locations that are likely to form the overexpressed regions we are interested in detecting. The symmetry of this distribution is believed to help in the detection. Furthermore, since the clusters of interest are definitely a minority of the probes, we can use this information to elicit informative prior distributions. For example, a prior that concentrates the probabilities $\{q_{0,K+1}, q_{1,K+1}, \ldots, q_{K,K+1}\}$ around small values. The detection procedure should be performed based on posterior statistics of the Z_i's. We propose a possible approach in Section 4.2.

4 Bayesian inference

Under the Bayesian paradigm, our aim is to obtain the posterior distribution of all unknown quantities involved in the model defined in the previous section. Given the high dimensionality and complexity of this posterior, we devise an MCMC algorithm to sample from it and then use Monte Carlo methods to perform estimation based on this sample.

4.1 The MCMC algorithm

The proposed MCMC is a Gibbs sampling with some Metropolis-Hastings (MH) steps. The blocking and sampling schemes are chosen in a way to favor fast convergence of the chain. In this direction, we introduce a set of independent auxiliary variables $V := \{V_i, \; i = 2,\ldots,n\}$ which allows direct sampling from the full conditional distribution of β (see Albert, 1992). Let $\vec{d_i} = (1, d_i)'$, we set

$$(V_i|\beta) \sim N(\beta'\vec{d_i}, 1) \quad \text{and} \quad W_i = \begin{cases} 1, & \text{if } V_i > 0\\ 0, & \text{if } V_i \leq 0. \end{cases}$$

Note that the original model for X is preserved – simply integrate V out to check it.

Let $\psi := \{\theta, \mu, \sigma^2\}$, we choose the following blocking scheme for the Gibbs sampler:

$$(Z, W) \ ; \ V \ ; \ (q_0, Q, \beta, \psi) \ ; \ \eta.$$

This scheme leads to a non-irreducible Markov chain which, in turn, does not converge. We circumvent this problem by making the algorithm a collapsed Gibbs sampling where V is integrated out from the full conditional distribution of (Z, W). This strategy makes the chain irreducible and guarantees its convergence to the target distribution (see Liu, 1994).

All the full conditional densities from the Gibbs sampler are proportional to the joint density of X and all the unknown quantities of the model – this density is given in (B.1), in Appendix B. We now describe each step of the algorithm.

Sampling V

The full conditional distribution of V is such that the V_i's are all independent with distribution $N(\beta' \vec{d}_i, 1)$, truncated to be positive if $W_i = 1$ and non-positive if $W_i = 0$.

Sampling (q_0, Q, β, ψ)

The four components of this block – q_0, Q, β and ψ, are conditionally independent. This means that a draw from this block is obtained by sampling each of the four components individually from their respective marginal full conditional distribution. The marginals of the first three components are given in (B.2)-(B.4), in Appendix B.

The marginal full conditional distribution of ψ is a truncation of a tractable and easy to simulate distribution, which is defined in (B.5). The truncation region is defined by the restriction $\{\theta_1 < \theta_2 < \ldots < \theta_K < \mu\}$. We sample exactly from this truncated distribution via rejection sampling by proposing from its non-truncated version and accepting if the restriction is satisfied. Since K is chosen to be small, this algorithm is computationally efficient.

Sampling η

The η_k's are sampled (jointly or individually) via Gaussian random walk MH step(s) properly tuned to have reasonable acceptance rates (see Roberts et al., 1997). Details on the proposal distribution and acceptance probability of the MH step where each η_k is sampled separately are given in (B.6) and (B.7), respectively, in Appendix B.

Sampling (Z, W)

This is the most challenging step of the MCMC. Although the full conditional can be factorised into multinomial distributions, obtaining such factorization and the parameters of each multinomial is not straightforward. Firstly, note that integrating out V provides the following full conditional kernel:

$$\begin{aligned} \pi(Z, W|\cdot) \quad \propto \quad & \prod_{i=1}^{n} \left[\prod_{k=1}^{K+1} \left[(f_k(X_i|\psi) q_{0k})^{Z_{ik}} \right]^{1-W_i} \left[(f_k(X_i|\psi) q_{k_{(i-1)}k})^{Z_{ik}} \right]^{W_i} \right] \\ & \times \left(\Phi_i^+\right)^{W_i} \left(\Phi_i^-\right)^{1-W_i} \prod_{k=1}^{K+1} q_{0k}^{Z_{0k}}, \end{aligned} \tag{8}$$

with $\Phi_i^+ = \Phi(\beta' \vec{d}_i)$ and $\Phi_i^- = \Phi(-\beta' \vec{d}_i)$.

Direct sampling from (8) is only possible if we adopt the following factorization:

$$\pi(Z, W|\cdot) \propto \pi(Z_1|\cdot) \prod_{i=2}^{n} \pi(Z_i|W_i, Z_{i-1}, \cdot) \pi(W_i|Z_{i-1}, \cdot),$$

which means that (Z, W) is sampled in the following order: $Z_1, W_2, Z_2, \ldots, W_n, Z_n$, i.e. forward in time. Each marginal distribution is multinomial (Bernoulli for the W_i's) for which the parameter values are obtained recursively and backward in time. Therefore, we denote this sampling scheme as a backward-filtering-forward-sampling (BFFS). Details on this sampling step are given in Appendix B.

The MCMC algorithm described here was implemented using the `R` programming language (R Core Team, 2016).

4.2 Cluster detection

As it was mentioned before in Section 3, the clusters of interest are expected to be accommodated by the Gaussian component of the mixture. In this sense, the cluster detection should be performed based on the posterior probability of one or more probes belonging to the Gaussian component of the mixture. Given a (approximate) sample of size M from the posterior distribution of Z and a sequence $\{i_1, i_2, \ldots, i_S\}$ of expressions ($S = 1, 2, \ldots$), the posterior probability of this sequence being a cluster of interest is given by:

$$\frac{1}{M} \sum_{m=1}^{M} \prod_{s=1}^{S} \mathbb{1}(Z_{i_s,K+1}^{(m)} = 1), \tag{9}$$

where $\mathbb{1}(.)$ is an indicator function.

A practical (computationally efficient) procedure for cluster detection is to evaluate the probability in (9) for each individual probe and look for sequences of probes of a minimum size (say 4 or 5) for which the individual probabilities are high (say > 0.5). For such sequences, the posterior probability in (9) will be considerably high (given the Markov dependence imposed by the model). One may also be interested in evaluating this probability for particular sequences – given some practical reason. In Section 5, we search for a sequence (size 5) of high individual posterior probabilities and report its corresponding joint posterior probability for each cancer data set.

5 Results for the cancer data sets

We consider five different sets of microarray data representing three types of cancer (see Section 2); the number of samples varies from 59 to 286 arrays. Since the expression scale is the same for every chromosome, we perform a joint analysis of all the 24 chromosomes (for each of the five data sets). Moreover, given that there is no reason to believe that there exists a dependence between the last and first probes of consecutive chromosomes (there is no distance to be measured), we fix the W_i corresponding to the first probe of each chromosome to be 0. An important implication of this is that the BFFS step to sample (Z, W) may be performed separately for each chromosome, given its conditional independence property. This leads to a significant improvement in the computational cost.

An important decision to fit the model is to choose the number K of gamma components in the mixture distribution. The histogram in Figure 1 (a) clearly indicates that a single gamma distribution would not be appropriate for this analysis due to the multimodal and asymmetric shape of the graph. In order to investigate other reasonable possibilities, we fit the model with $K = 2$, 3, 4 and 5 to the "Breast 1" data set (the distributions of the expressions have similar behavior for all data sets). Results are reported in Appendix D and indicate that $K = 4$ is the best choice. The model with $K = 2$ led to a wide (high variance) Gaussian component including too many expressions (some of them very low). Results for $K = 3$ are similar but still significantly different from those with $K = 4$ which, in turn, are virtually the same as for $K = 5$ (especially in terms of overexpressed region detection).

In terms of prior specification, consider again the notation in (7). We assume $r_0 = (750, 750, 750, 750, 10)'$ as the parameter vector of the Dirichlet prior for q_0; here, we indicate that an expression at a location without a Markov dependence most likely belongs to a gamma component in the mixture. The large equal weights for the gammas are required in this application involving thousands of locations; small concentration parameters would lead to a weakly informative prior dominated by the data. This strategy is also considered for the prior specification of q_k. We assume the following matrix:

$$r = \begin{pmatrix} 969.70 & 484.85 & 96.97 & 48.48 & 10.0 \\ 484.85 & 969.70 & 484.85 & 96.97 & 10.0 \\ 96.97 & 484.85 & 969.70 & 484.85 & 10.0 \\ 48.48 & 96.97 & 484.85 & 969.70 & 10.0 \\ 48.48 & 48.48 & 96.97 & 484.85 & 969.7 \end{pmatrix}.$$

The actual values of the hyperparameters for q_0 and Q are established in a way that the prior had the same weight for the Gaussian component and the same level of information for each specification of K in the above mentioned sensitivity analysis (rounded numbers were fixed for $K = 3$). The k-th row of matrix r is the parameter vector of the Dirichlet prior for q_k; recall that here we account for the Markov dependence. Note

that the largest weights were defined in the main diagonal to favor the model to allocate the i-th observation in the same component of its neighbour from location $i-1$. This matrix also has weights decreasing as we move away from the main diagonal elements – this configuration is used to discourage the model to assign two consecutive observations in distant components. Another important aspect of this prior are the small weights specified for the Gaussian component in the last column of the first four rows and the high weight specified in the last row; this choice is an strategy to make this component accessible only to those observations having great evidence of overexpression and having neighbours with the same characteristics.

The bivariate normal distribution for the coefficients (β_0, β_1) is set with mean (4,−8) and covariance matrix $10\,\boldsymbol{I_2}$; i.e., we assume prior independence for these coefficients and the mean indicates the following information about the probability of having a Markov dependence: $\lim_{d_i \to 0} \rho_i \approx 1$, $\lim_{d_i \to 1} \rho_i \approx 0$ and $\rho_i = 0.5$ for $d_i = 0.5$.

In this application, we assume a gamma prior distribution with mean 50 and shape 1 for the shape parameter η_k of all the four gamma components in the mixture; its variance is 50^2. In addition, the mean θ_k has the following inverse gamma prior specification: $t_{1k} = 4$ (for $k = 1, 2, 3, 4$), $t_{21} = 9$, $t_{22} = 18$, $t_{23} = 27$ and $t_{24} = 36$. These hyperparameter choices provide the expected values (3, 6, 9, 12) and the standard deviations (2.12, 4.24, 6.36, 8.48) for $k = 1, 2, 3, 4$, respectively. We have chosen the expected values based on the scale of the observations and the role we expect the gamma components to play in the model. Finally, we specify a normal-inverse-gamma prior with parameters (15, 25, 2.1, 1.1) for the pair (μ, σ^2) related to the Gaussian component. This implies that $E(\sigma^2) = 1$ and $\text{Var}(\sigma^2) = 10$. A prior sensitivity analysis was performed by doubling the standard deviations of $(\mu, \sigma^2, \theta_1, \ldots, \theta_4)$ and the very same results were obtained – see Appendix D.

Table 1: Posterior mean and standard deviation (in parentheses) for the parameters of the mixture distribution with four gamma and one Gaussian components.

	Breast 1	*Breast 2*	*Breast 3*	*Ovarian*	*Brain*
θ_1	4.020 (0.008)	3.769 (0.010)	4.127 (0.008)	3.642 (0.007)	4.365 (0.007)
η_1	105.775 (2.920)	112.609 (4.107)	120.475 (3.386)	177.937 (5.546)	268.408 (8.919)
θ_2	6.398 (0.028)	5.789 (0.051)	6.351 (0.029)	4.991 (0.029)	5.612 (0.032)
η_2	20.815 (0.326)	22.487 (0.573)	22.724 (0.399)	36.937 (0.995)	52.620 (1.721)
θ_3	8.066 (0.021)	7.684 (0.021)	8.080 (0.023)	7.213 (0.025)	7.776 (0.030)
η_3	62.996 (1.596)	65.062 (1.500)	57.345 (1.556)	51.707 (1.238)	61.648 (1.546)
θ_4	9.824 (0.028)	9.473 (0.033)	9.882 (0.027)	8.846 (0.029)	9.573 (0.031)
η_4	94.436 (2.528)	87.524 (2.750)	93.658 (2.547)	68.373 (1.982)	88.486 (2.623)
μ	12.132 (0.033)	11.997 (0.036)	12.261 (0.032)	11.504 (0.034)	11.780 (0.034)
σ^2	0.831 (0.034)	0.947 (0.040)	0.843 (0.033)	1.103 (0.039)	0.873 (0.033)

Table 2: Posterior mean of the coordinates of the Z_i's – posterior weights of the mixture components.

k	*Breast 1*	*Breast 2*	*Breast 3*	*Ovarian*	*Brain*
1	0.138	0.150	0.140	0.120	0.124
2	0.387	0.378	0.381	0.278	0.269
3	0.274	0.286	0.280	0.371	0.365
4	0.165	0.148	0.162	0.188	0.198
5	0.036	0.037	0.037	0.042	0.045

Table 3: Comparison involving pairs of data sets. Number of probes identified in the Gaussian component in both data sets of the pair, relative to the number of probes identified in the Gaussian component in at least one of the data sets of the pair.

	Breast 1	*Breast 2*	*Breast 3*	*Ovarian*	*Brain*
Breast 1	100%	79.1%	87.4%	71.6%	62.9%
Breast 2		100%	84.6%	76.1%	64.5%
Breast 3			100%	75.4%	65.9%
Ovarian				100%	69.3%
Brain					100%

The MCMC is set to perform 15,000 iterations with a burn-in of 5,000. Let $\mathbf{1_{(l_1 \times l_2)}}$ be a $l_1 \times l_2$ matrix of ones; in terms of initial values, we consider $q_0^{(0)} = 0.2\,\mathbf{1_{(5\times 1)}}$ and $Q^{(0)} = 0.2\,\mathbf{1_{(5\times 5)}}$. In addition, we set $\beta_0^{(0)} = 4$, $\beta_1^{(0)} = -8$ as specified in their prior distribution. The initial values of the remaining parameters are set based on statistics of the data; we break the support of the histogram (see Figure 1) into 4 contiguous intervals and then evaluate their means and variances to determine: $\theta_1^{(0)} = 4.43$, $\theta_2^{(0)} = 6.70$, $\theta_3^{(0)} = 8.96$, $\theta_4^{(0)} = 11.23$, $\mu^{(0)} = 13.50$, $\eta_1^{(0)} = 50.87$, $\eta_2^{(0)} = 108.30$, $\eta_3^{(0)} = 200.30$, $\eta_4^{(0)} = 330.93$ and $(\sigma^2)^{(0)} = 0.21$. The variances of the random walk proposals in the MH step to sample η_1, η_2, η_3 and η_4 were chosen to provide reasonable (around 44%) acceptance rates. Convergence is rapidly attained – Appendix E shows some diagnostics.

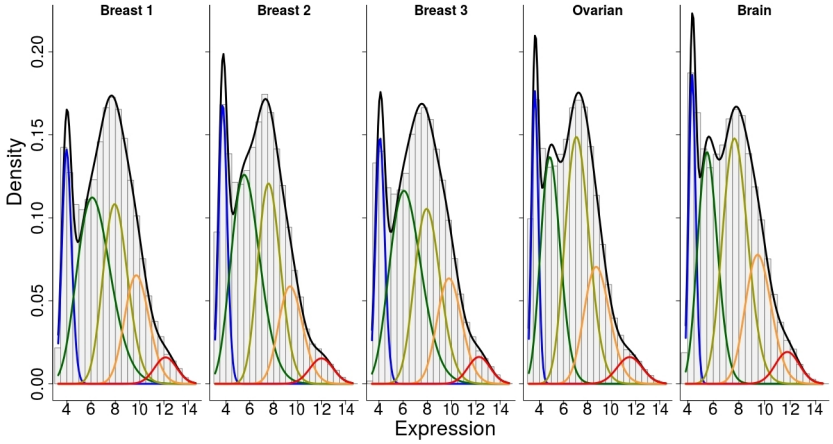

Figure 3: Histogram of all expressions overlaid by the estimated mixture density (black curve) and its components (gammas in blue, green, yellow and orange, normal in red).

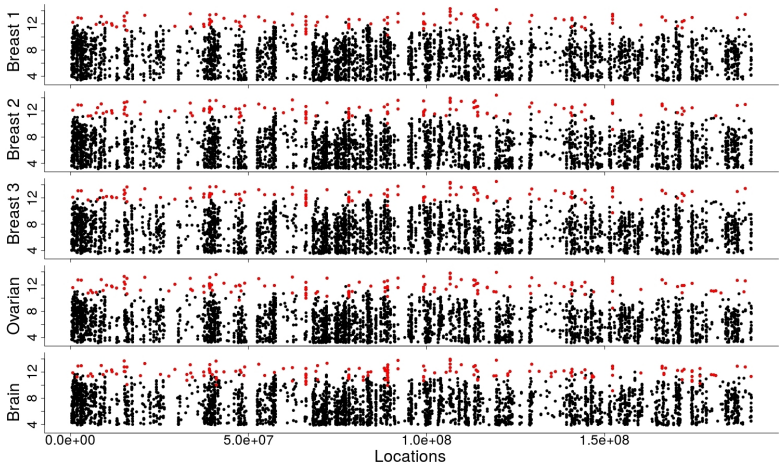

Figure 4: Expression values along chromosome 4. The red dots indicate observations where the Gaussian component has posterior probability above 0.5.

Table 1 presents the posterior estimates for each component of the mixture distribution and Table 2 their corresponding posterior weight (posterior mean of the coordinates of the Z_i's). For each parameter, note that the means and standard deviations are similar across the data sets; this suggests mixture densities having nearly the same shape when comparing different cancers.

Figure 3 shows the histogram of the data overlaid by the estimated mixture density and their components for each data set. Note that the range of the Gaussian is not too wide and, therefore, accommodates the expressions in the right tail of the histograms. This result suggests that the mixture with $K = 4$ gamma components represents well the data.

Figure 4 shows the behaviour of the expression values along chromosome 4 for all five data sets; here, we explore a single chromosome since the graph displaying all of them would be difficult to visualise. As

can be seen, the red dots, representing observations with high posterior probability (> 0.5) of belonging to the Gaussian component, are concentrated on the top of the graphs. Note that non-Gaussian observations (black dots) can be identified among Gaussian ones, which indicates high expressions originated from the right tail of a gamma distribution in the mixture. This occurs as a result of the model structure accounting for distances between locations to determine the strength of the Markov dependence. Again, the simple strategy of choosing a threshold above which all points are in red would inflate the number of overexpressed regions.

We also compute the posterior probability of a specific cluster of 5 locations being an overexpressed region for each data set – see (9). We evaluate the sequence composed by locations 33,312 to 33,316. Its posterior probability of being an overexpressed cluster is: 0.9996 (Breast 1), 0.9998 (Breast 2), 0.9998 (Breast 3), 0.9998 (Ovarian) and 0.9989 (Brain).

Another interesting aspect displayed in Figure 4 is the fact that for the same location the expressions from each data set are similar, consequently, the cluster identification tends also to be similar comparing these data sets. If we consider all 24 chromosomes (91,090 locations identified via BLAT), the number of probes in the Gaussian component are: 3,132 (Breast 1), 3,277 (Breast 2), 3,271 (Breast 3), 3,723 (Ovarian) and 3,860 (Brain). Table 3 compares each pair of data sets and indicates the frequency of locations identified in the Gaussian component in both data sets, relative to the total number of locations identified in the Gaussian component in at least one of them. Note that the largest percentages (79.1%–87.4%) are obtained for pairs involving two breast cancer data sets; this is a plausible result suggesting that the model is working well in this cluster identification problem involving 91,090 locations. The overexpressed regions in the breast cancer data sets seem to have higher similarity with the ovarian cancer (71.6%–76.1%) than with the brain cancer (62.9%–69.3%).

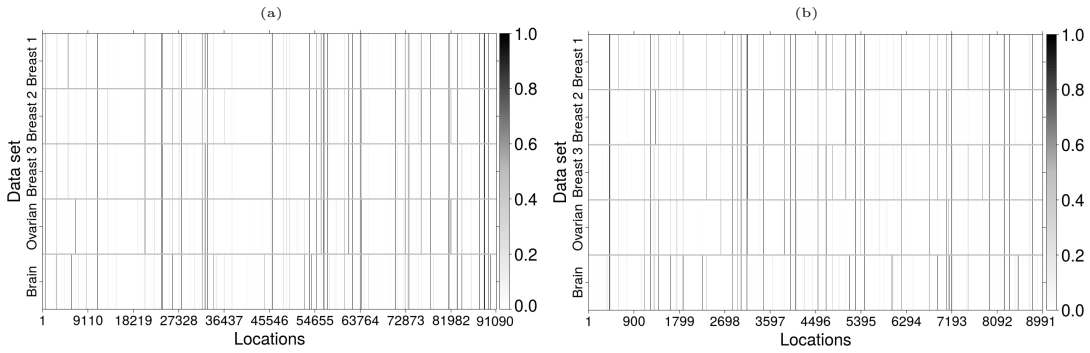

Figure 5: Heat map image indicating for each location the posterior probability of belonging to the Gaussian component. Panel (a) represents all chromosomes and Panel (b) shows chromosome 1.

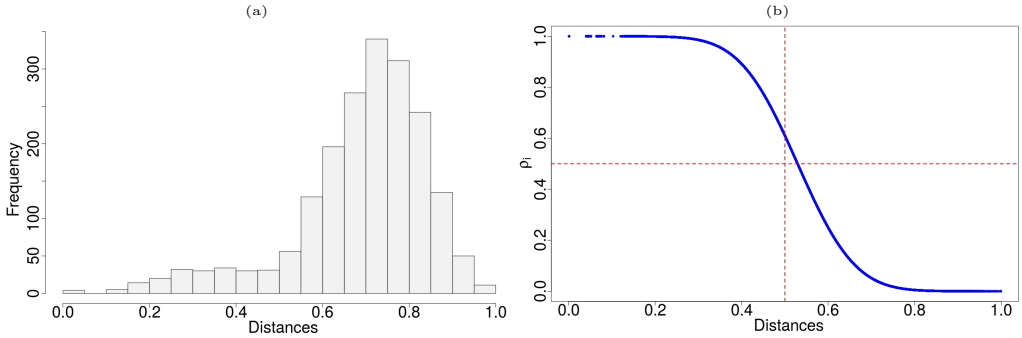

Figure 6: Panel (a): Histogram of the distances between the location of a cluster, with a single observation, and its nearest non-Gaussian neighbours. Panel (b): Estimated relationship between the distances and ρ_i (probability of having a Markov dependence with the previous probe); here, we consider the "Breast 1" data set and the posterior mean of β_0 (5.05) and β_1 (-9.55) to draw the curve.

Figure 5 shows a grey scale heat map investigating the posterior probabilities of being a Gaussian component observation. The results for all 24 chromosomes are shown in panel (a) and only the chromosome 1 is explored in panel (b); images for the chromosomes 2-22, X and Y have a similar behaviour and are

presented (individually) in Appendix C. The conclusions here reinforce the previous analysis of Figure 4 and Table 3; i.e., the data sets tend to agree, exhibiting similar probabilities for the same location. Note that a coherent pattern of grey/black probabilities can be clearly seen in both panels and the majority of locations are in white due to a estimated probability of Gaussian classification near 0 for most chromosomal regions. This coherence between results from different microarray data sets, specially between those three for breast cancer, is a strong evidence that the proposed model is performing well. A disagreement between data sets for the same type of tumor is not expected and, assuming that the data were correctly preprocessed and aligned, this would suggest a critical model issue; this is not the case here.

In this application, we have observed different sizes of sequences of consecutive locations belonging to the Gaussian component when adopting the classification rule described in the caption of Figure 4. The shortest case is a single observation surrounded by two non-Gaussian neighbours; the longest sequence, comparing all data sets, involves 17 locations. One may consider that the identification of a single Gaussian location with non-Gaussian neighbours is potentially a model issue; however, the histogram in Figure 6 (a) shows that this is not the case. It shows that the distribution of all distances between a single Gaussian location and its non-Gaussian neighbours is highly concentrated in the region above 0.5. Therefore, the magnitude of these isolated expressions is compatible with an overexpressed cluster and the model cannot use the neighbourhood information to change this classification because the neighbours are too far away. These locations belong to potential clusters but a confirmation relies on their non-observed neighbours.

Figure 6 (b) presents the estimated relationship between the distances and the probability ρ_i of having a Markov dependence at location i for the "Breast 1" data; two dashed lines were included to highlight the probability and the distance 0.5. The posterior means of (β_0,β_1) for each data set are: (5.05,-9.55) in "Breast 1", (5.02,-9.50) in "Breast 2", (5.26,-9.85) in "Breast 3", (5.27,-9.83) in "Ovarian" and (4.92,-9.31) in the "Brain" data set. These are all similar, which implies that their corresponding curves will be very close to the one in panel (b). We can evaluate the decreasing behaviour of the curve to understand the impact of the distances on ρ_i; for example, the slow decay before the distance 0.4 indicates that the Markov dependence is strong for those distances. In addition, the curve has a fast decay between distances 0.4 and 0.6, and is approximately 0 from 0.8 onwards.

6 Conclusions

In this paper, we have developed a hidden Markov model designed for an application involving Affymetrix DNA microarray data. The study focuses on the identification of chromosomal regions associated with high gene expression measurements, which we call overexpressed regions. The microarray probe expressions are mapped to locations in the human chromosomes using the alignment algorithm BLAT. As a result, the data is configured as an irregularly spaced sequence along the chromosomes; five different data sets representing breast, ovarian and brain cancers were considered in the analysis. The original light intensities from the arrays were preprocessed via RMA.

The proposed model assumes a mixture distribution with four gamma and one normal components to cluster the observations. The largest mean is imposed for the Gaussian component, which is supposed to accommodate the target overexpressed values. The model takes advantage of the distances between the identified locations to determine whether there is a Markov dependence between neighbours and uses such dependence to (stochastically) define the target clusters. Inference is performed under the Bayesian paradigm via an MCMC algorithm consisting of a carefully devised Gibbs sampling.

The results indicate that the model is selective in the cluster determination, being able to discriminate well the locations to identify the overexpressed regions; on average, only 3.8% of the locations are classified (for a suitable classification rule) in potential overexpressed clusters. Looking at the intersection of overexpressed regions from the analised data sets and evaluating the coherence pattern exhibited in image graphs, our findings suggest great similarity between them. As expected, the greatest similarity is observed among the breast cancer data sets, which is a strong indication that the model is performing well. The real application also shows that the breast cancer data is more similar to the ovarian cancer than to the brain cancer.

We perform a general cluster identification analysis to present a global picture of the overexpressed regions across the human genome. This means that the results obtained here can be used for a more detailed description of specific regions of interest in any chromosome; for example, one can easily compute

the probability of a particular set of probes being an overexpressed cluster through the cluster detection procedure presented in (9).

Acknowledgements

The authors would like to thank the Associate Editor and the anonymous referee for their constructive comments leading to a much improved version of the paper. The authors also thank FAPEMIG for supporting this research. The first author would like to thank Joseph E. Lucas for the initial discussion about this topic at Duke University.

Appendix A: Details about BLAT and RMA

The BLAT algorithm was developed to quickly find genome positions having at least 95% similarity with sequences of interest in a database. In fact, BLAT is available for different types of query sequences and, in our application, we use this algorithm on DNA to search for chromosome locations having high compatibility with the 25-bases probe sequences from the microarray. The shorter is the length of the sequence, the higher is the risk of incorrect mapping; this issue is not peculiar to BLAT and it can also occur in other alignment techniques. However, as indicated in Allen et al. (2011), BLAT is a popular algorithm being used in several researches to align probe sequences to a recent release of the human genome or transcriptome. The main advantage of the method is its low computational cost; BLAT does not keep the whole genome information in memory allowing high performance in an ordinary Linux computer. In addition, BLAT can be accessed for small searches through a web server at `https://genome.ucsc.edu/cgi-bin/hgBlat`. The sources and executables to install and run large mapping jobs in a computer are freely available for academic, personal and nonprofit purposes; see `http://www.kentinformatics.com` for details.

The results obtained via BLAT may contain some inconsistencies; a probe may be aligned to multiple locations and a location may be associated with two or more probes. We have no additional information to correct this mapping problem, which affects very few probes and locations; therefore, the involved expression values will be removed from the study without compromising the analysis.

In the data preprocessing routine, the intensity values are transformed using the standard $\log_2$-base scale and these measurements are adjusted within each chip and across the replicate arrays. Different sources can cause distortions in the data, for example, cross-hybridization, dust, chip defect, the amount of RNA in the samples, camera exposure time, scanner calibration, etc. The first three sources affect observations within a single chip and the remaining ones introduce noise between chips.

Among the different techniques to preprocess the data, we consider the Robust Multi-chip Average (RMA) (see Irizarry et al., 2003). It has three main steps: background correction, quantile normalization and summarization. In short, the first step fits a linear model "signal plus error" to explain the intensity of probe i in the array j; here, the main goal is to estimate the signal component. The quantile normalization step adjusts the probe intensities to ensure that measurements from different arrays are comparable. Finally, the summarization step calculates a single value representing the expressions of the probes in each probe set.

In our study, we apply only the first two RMA preprocessing steps. The third step is ignored because BLAT aligns the probe level data and not the probe set summarised measurement. In order to apply the RMA background correction and quantile normalization, we use the software `R` (R Core Team, 2016) and its package `affy` (Gautier et al., 2004) integrated into the collaborative project Bioconductor (`http://www.bioconductor.org`) providing tools for computational biology (Gentleman et al., 2004). We also highlight the fact that both RMA and BLAT work only with PM probes; therefore, the MM probes are not considered in our final aligned data.

Appendix B: Joint and full conditional distributions

The joint density of X and all the unknown quantities of the model is given by:

$$\begin{aligned}
& \pi(X, Z, W, V, q_0, Q, \psi, \beta) = \\
= & \left[\prod_{i=1}^{n} \pi(X_i|Z_i, \psi)\pi(Z_i|Z_{i-1}, W_i, q_0, Q)\pi(W_i|V_i)\pi(V_i|\beta)\right] \pi(q_0, Q, \psi, \beta) \\
= & \left[\prod_{i=1}^{n}\left[\prod_{k=1}^{K+1} f_k(X_i|\psi)^{Z_{i,k}} (q_{0k}^{Z_{i,k}})^{1-W_i} (q_{k_{(i-1)}k}^{Z_{i,k}})^{W_i}\right]\right. \\
& \times \left[\mathbb{1}(W_i = 1)\, \mathbb{1}(V_i > 0) + \mathbb{1}(W_i = 0)\, \mathbb{1}(V_i \leq 0)\right] \\
& \left.\times\, \phi(V_i - \beta' \vec{d}_i)\right] \left[\prod_{k=1}^{K+1} q_{0k}^{Z_{0,k}} q_{0k}^{r_{0k}-1}\right] \left[\prod_{k_1=1}^{K+1}\prod_{k=1}^{K+1} q_{k_1 k}^{r_{k_1 k}-1}\right] \pi(\psi)\, \pi(\beta),
\end{aligned} \tag{B.1}$$

where $k_{(i-1)} = j$ if $Z_{i-1,j} = 1$, $\pi(\beta) = |\Sigma_0|^{-1/2}\phi_2[\Sigma_0^{-1/2}(\beta - \mu_0)]$ and

$$\begin{aligned}
\pi(\psi) = & \left[\prod_{k=1}^{K} \pi_{IG}(\theta_k; t_{1k}, t_{2k})\, \mathbb{1}(\theta_1 < \ldots < \theta_K)\, \pi_G(\eta_k; e_{1k}, e_{2k})\right] \\
& \times \left[\pi_{NG}(\mu, \sigma^2; m, v, s_1, s_2)\, \mathbb{1}(\mu > \max_{1 \leq k \leq K}\{\theta_k\})\right].
\end{aligned}$$

The full conditional distribution of (q_0, Q, ψ, β) is:

$$(q_0|\cdot) \sim \text{Dir}\left[r_0 + \sum_{i=1}^{n} Z_i(1 - W_i)\right]; \tag{B.2}$$

$$(q_k|\cdot) \sim \text{Dir}\left[r_k + \sum_{i=2}^{n} (Z_{i-1,k} W_i) Z_i\right]; \tag{B.3}$$

$$(\beta|\cdot) \sim N_2(\mu_0^*, \Sigma_0^*), \tag{B.4}$$

with $\Sigma_0^* = \left(\Sigma_0^{-1} + \sum_{i=1}^{n} \vec{d}_i \vec{d}_i{}'\right)^{-1}$ and $\mu_0^* = \Sigma_0^* \left(\Sigma_0^{-1}\mu_0 + \sum_{i=1}^{n} V_i \vec{d}_i\right)$.

$$(\sigma^2|\cdot) \sim IG(s_1^*, s_2^*); \quad (\mu|\cdot) \sim N(m^*, v^*); \quad (\theta_k|\cdot) \sim IG(t_{1k}^*, t_{2k}^*),\ k = 1, \ldots, K; \tag{B.5}$$

where $s_1^* = s_1 + \frac{1}{2}\left(\sum_{i=1}^{n} Z_{i,K+1}\right)$ and
$s_2^* = s_2 + \frac{1}{2}\left[\frac{m^2}{v} + \sum_{i=1}^{n} Z_{i,K+1} X_i^2 - \frac{v}{1 + v\sum_{i=1}^{n} Z_{i,K+1}}\left(\frac{m}{v} + \sum_{i=1}^{n} Z_{i,K+1} X_i\right)^2\right]$;
$v^* = \frac{v\sigma^2}{1 + v\sum_{i=1}^{n} Z_{i,K+1}}$ and $m^* = v^*\left(\frac{m + v\sum_{i=1}^{n} Z_{i,K+1} X_i}{v\sigma^2}\right)$;
$t_{1k}^* = t_{1k} + \eta_k \sum_{i=1}^{n} Z_{i,k}$ and $t_{2k}^* = t_{2k} + \eta_k \sum_{i=1}^{n} Z_{i,k} X_i$.

The MH algorithm to sample η_k takes into account:

- Proposal distribution:

$$\eta_k^* \sim N(\eta_k, \tau^2), \text{ where } \eta_k \text{ is the current value of the chain.} \tag{B.6}$$

- Acceptance probability:

$$\min\left\{1\ ,\ \frac{\prod_{i;\ Z_{i,k}=1} f_k(X_i|\eta_k^*, \theta_k)\pi(\eta_k^*)}{\prod_{i;\ Z_{i,k}=1} f_k(X_i|\eta_k, \theta_k)\pi(\eta_k)}\right\}; \tag{B.7}$$

where θ_k is the current value of parameter θ_k.

The sampling step of (Z, W) is as follows:

$$\begin{aligned} Z_1 &\sim \text{Mult}(1, q_0^*); \\ (W_i | Z_{i-1,j} = 1, \cdot) &\sim \text{Ber}(1, p^*_{(i,j)}),\ i = 2, \ldots, n; \\ (Z_i | W_i = l, Z_{i-1,j} = 1, \cdot) &\sim \text{Mult}(1, q^*_{(i,j,l)}),\ i = 2, \ldots, n \end{aligned}$$

where

$$\begin{aligned} q^*_{0k} &= c_{1,k} q_{0k} / a_1; \\ p^*_{(i,j)} &= b_{i,j} \Phi_i^+ / c_{i,j}; \\ q^*_{(i,j,0)k} &= c_{i+1,k} f_k(X_i|\psi) q_{0k} / a_i; \quad q^*_{(i,j,1)k} = c_{i+1,k} f_k(X_i|\psi) q_{j,k} / b_{i,j}. \end{aligned} \tag{B.8}$$

Consider $c_{n+1,j} = 1, \ \forall j, \ $ and:

$$\begin{aligned} a_n &= \sum_{k=1}^{K+1} f_k(X_n|\psi) q_{0k}; \quad b_{n,j} = \sum_{k=1}^{K+1} f_k(X_n|\psi) q_{j,k}; \quad c_{n,j} = b_{n,j} \Phi_n^+ + a_n \Phi_n^-; \\ a_i &= \sum_{k=1}^{K+1} c_{i+1,k} f_k(X_i|\psi) q_{0k}; \ \ b_{i,j} = \sum_{k=1}^{K+1} c_{i+1,k} f_k(X_i|\psi) q_{j,k}; \ \ c_{i,j} = b_{i,j} \Phi_i^+ + a_i \Phi_i^-, \\ &\quad \text{for } i = n-1, \ldots, 2; \\ a_1 &= \sum_{k=1}^{K+1} c_{1,k} q_{0k}. \end{aligned} \tag{B.9}$$

In the filtering procedure, the calculations are performed recursively, starting from n and moving backwards. On each step, we compute the scalar a_i and the $(K+1)$-vectors b_i and c_i as indicated in (B.9). Next, we compute the probabilities in (B.8) to update the auxiliary variables Z_i and W_i.

Appendix C: Additional results of the real data application

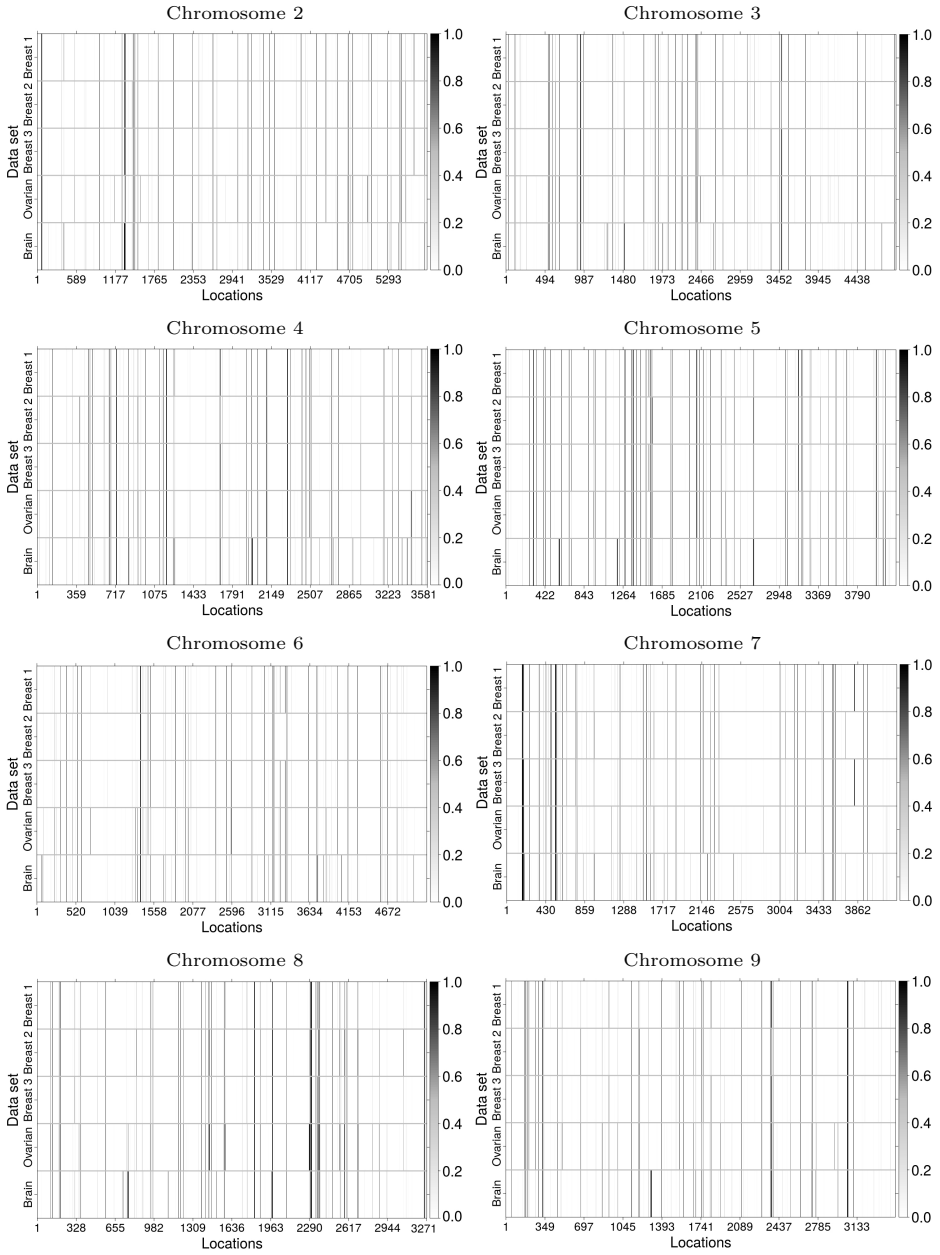

Figure C.1: Heat map image (chromosomes 2-9) indicating for each location the posterior probability of belonging to the Gaussian component.

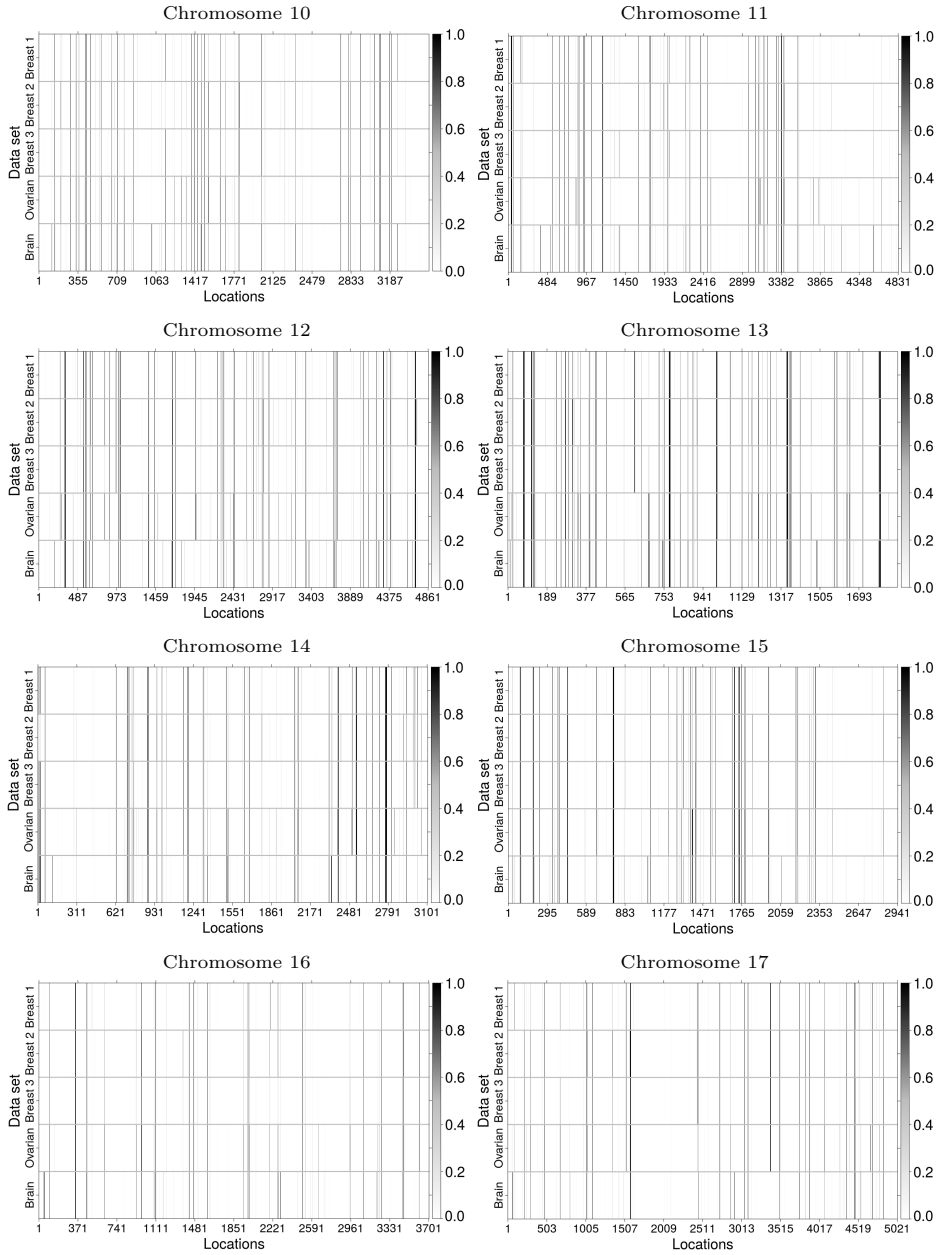

Figure C.2: Heat map image (chromosomes 10-17) indicating for each location the posterior probability of belonging to the Gaussian component.

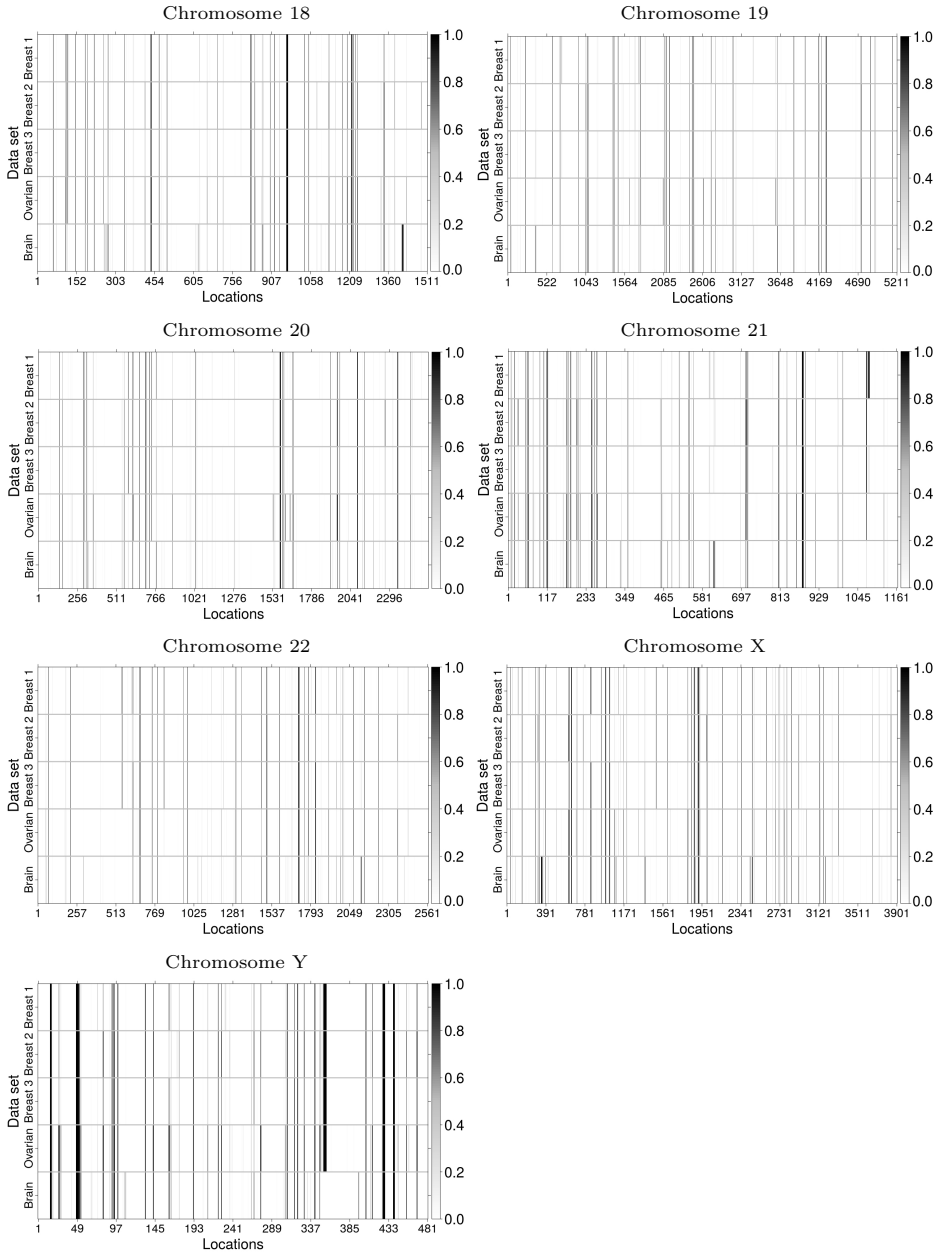

Figure C.3: Heat map image (chromosomes 18-22, X and Y) indicating for each location the posterior probability of belonging to the Gaussian component.

Appendix D: Sensitivity analysis

In order to investigate the sensitivity of the results to the choice of K (number of gamma components), we fit the model under four different specifications: $K = 2$, 3, 4 and 5. The Dirichlet priors for q_0 and Q are specified in a way that the Gaussian component's weight and the degree of information in the prior (sum of the hyperparameters of the Dirichlet) are the same across all configurations. Results are presented in Figure D.1 and Table D.1. They indicate that $K = 2$ does not provide good results as the Gaussian component's variance is too high to accommodate overexpressed probes. Results are quite similar (but significantly different) for $K = 3$ and $K = 4$, and they are practically the same for $K = 4$ and $K = 5$. In particular, two of the gamma components for $K = 5$ are virtually the same – the posterior mean of (θ, η) for each of these components are $(7.81, 49.23)$ and $(7.90, 46.89)$. Hence, we choose the model with $K = 4$,

for which results are reported in Section 5.

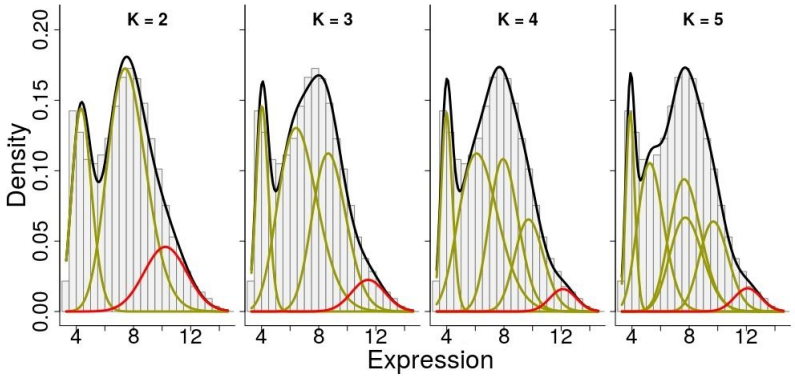

Figure D.1: Histogram of all expressions ("Breast 1" data) overlaid by the estimated mixture density (black curve) and its components (gammas in yellow, normal in red), for all the four values of K.

Table D.1: Posterior estimates (for each K, "Breast 1" data) of the weight, μ and σ^2 related to the Gaussian component, and the mean (Expec.) and variance (Var.) of the (normalised) mixture of gammas; standard errors in parentheses.

K	weight	μ	σ^2	Expec.	Var.
2	0.1715 (0.0033)	10.2320 (0.0281)	2.2208 (0.0362)	6.6956 (0.0103)	2.3649 (0.0177)
3	0.0654 (0.0017)	11.4605 (0.0341)	1.3443 (0.0388)	7.0111 (0.0074)	2.8614 (0.0209)
4	0.0365 (0.0011)	12.1317 (0.0330)	0.8314 (0.0338)	7.1193 (0.0062)	3.2746 (0.0226)
5	0.0383 (0.0012)	12.0889 (0.0354)	0.8599 (0.0354)	7.1112 (0.0060)	3.4991 (0.0204)

Table D.2: Comparison of prior specifications ("Breast 1" data). Posterior estimates of the weight, μ and σ^2 related to the Gaussian component, and the mean (Expec.) and variance (Var.) of the (normalised) mixture of gammas; standard errors in parentheses.

Prior	weight	μ	σ^2	Expec.	Var.
Original	0.0365 (0.0011)	12.1317 (0.0330)	0.8314 (0.0338)	7.1193 (0.0059)	3.2746 (0.0226)
Vague	0.0364 (0.0011)	12.1338 (0.0333)	0.8305 (0.0334)	7.1197 (0.0059)	3.2736 (0.0223)

We also perform a prior sensitive analysis for the parameters of the mixture components, with $K = 4$. We consider two different prior specifications, where the standard deviations from one of them are twice the standard deviations from the other one. More specifically, we assume in specification 1: $(\mu, \sigma^2) \sim NIG(15,\ 25,\ 2.1,\ 1.1)$, $\theta_k \sim IG(4, t_{2k})$ with $t_{2k} = 9,\ 18,\ 27,\ 36$, for $k = 1, 2, 3, 4$, respectively; this assigns means (3, 6, 9, 12) and standard deviations (2.12, 4.24, 6.36, 8.48) for each of the four gamma components. Note that this configuration is the original one explored in Section 5. In specification 2, we set: $(\mu, \sigma^2) \sim NIG(15,\ 100,\ 2.025,\ 1.025)$, $\theta_k \sim IG(2.5, t_{2k})$ with $t_{2k} = 4.5,\ 9.0,\ 13.5,\ 18.0$, for $k = 1, 2, 3, 4$, respectively; this provides the same means, but larger standard deviations (4.24, 8.48, 12.72, 16.96) for each of the four gamma components. Results are shown in Figure D.2 and Table D.2. They are visually the same indicating robustness to the prior specification for these parameters.

Appendix E: MCMC diagnostics

As we have mentioned in the text, the MCMC we design has good convergence properties. This is due to the chosen blocking scheme and the fact that we can sample directly from the full conditional distribution of (Z, W). Figures E.1 and E.2 present some graphs of the MCMC chain for the block (Z, W) to support this point. The trace plots show good mixing properties and the ergodic average trajectories show fast convergence.

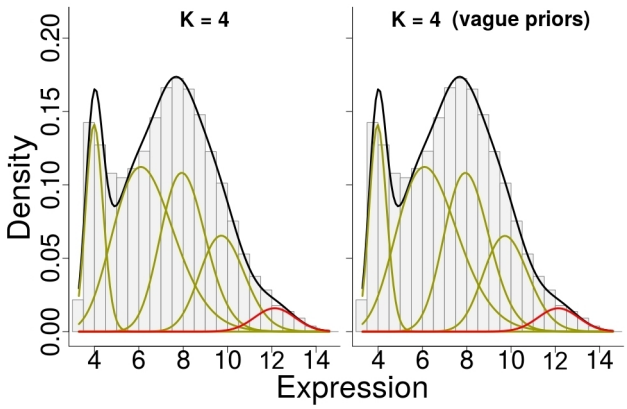

Figure D.2: Histogram of all expressions ("Breast 1" data) overlaid by the estimated mixture density (black curve) and its components (gammas in yellow, normal in red), for the two prior specifications.

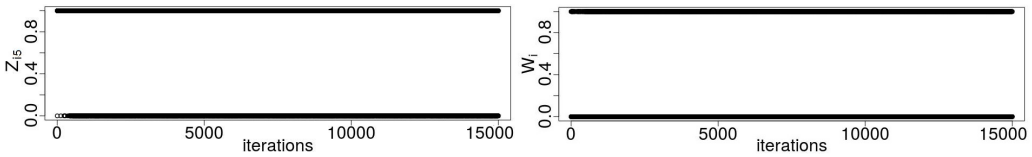

Figure E.1: Trace plots of a pair (Z_{i5}, W_i) with posterior mean around 0.5 ("Breast 1" data).

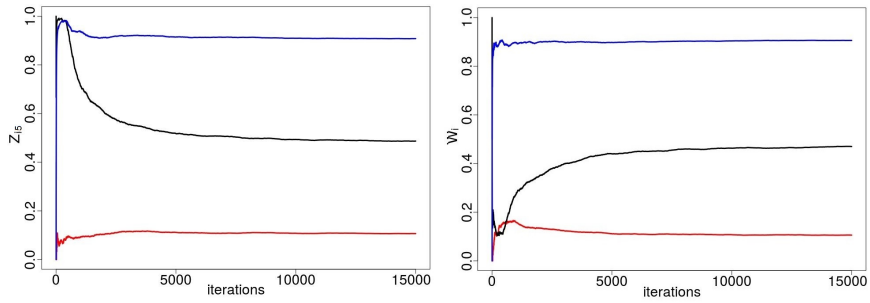

Figure E.2: Evolution of the ergodic average of three pairs (Z_{i5}, W_i) along the MCMC chain ("Breast 1" data).

References

Albert, J. H. (1992). Bayesian estimation of normal ogive item response curves using Gibbs sampling. *Journal of Educational and Behavioral Statistics, 17*, 251–269.

Allen, J. D., S. Wang, M. Chen, L. Girard, J. D. Minna, Y. Xie, and G. Xiao (2011). Probe mapping across multiple microarray platforms. *Briefings in Bioinformatics 13*(5), 547–554.

Amaratunga, D., J. Cabrera, and Z. Shkedy (2014). *Exploration and analysis of DNA microarray and other high-dimensional data* (2 ed.). Hoboken: Wiley.

Autio, R., S. Hautaniemi, P. Kauraniemi, O. Yli-Harja, J. Astola, M. Wolf, and A. Kallioniemi (2003). CGH-Plotter: MATLAB toolbox for CGH-data analysis. *Bioinformatics 22*, 1714–1715.

Barry, D. and J. A. Hartigan (1992). Product partition models for change point problems. *The Annals of Statistics 20*(1), 260–279.

Baudry, J. P., A. E. Raftery, G. Celeux, K. Lo, and R. Gottardo (2010). Combining mixture components for clustering. *Journal of Computational and Graphical Statistics 19*(2), 332–353.

Beskos, A., O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead (2006). Exact and computation efficient likelihood-based estimation for discretely observed diffusion processes (with discussion). *Journal of the Royal Statistical Society, Series B 68*(3), 333–382.

Bivand, R. and G. Piras (2015). Comparing implementations of estimation methods for spatial econometrics. *Journal of Statistical Software 63*(18), 1–36.

Broet, P., A. Lewin, S. Richardson, C. Dalmasso, and H. Magdelenat (2004). A mixture model-based strategy for selecting sets of genes in multiclass response microarray experiments. *Bioinformatics 20*, 2562–2571.

Broet, P. and S. Richardson (2006). Detection of gene copy number changes in CGH microarrays using a spatially correlated mixture model. *Bioinformatics 22*, 911–918.

Broet, P., S. Richardson, and F. Radvanyi (2002). Bayesian hierarchical model for identifying changes in gene expression from microarray experiments. *Journal of Computational Biology 9*, 671–683.

Carvalho, B. S. and R. A. Irizarry (2010). A framework for oligonucleotide microarray preprocessing. *Bioinformatics 26*(19), 2363–2367.

Cheung, V. G., R. S. Spielman, K. G. Ewens, T. M. Weber, M. Morley, and J. T. Burdick (2005). Mapping determinants of human gene expression by regional and genome-wide association. *Nature 437*(27), 1365–1369.

Chib, S. (1998). Estimation and comparison of multiple change-point models. *Journal of Econometrics 86*, 221–241.

Dalma-Weiszhausz, D. D., J. Warrington, E. Y. Tanimoto, and C. G. Miyada (2006). *Methods in Enzymology*, Volume 410 of *DNA Microarrays, Part A: Array Platforms and Wet-Bench Protocols*, Chapter The Affymetrix GeneChip platform: an overview, pp. 3–28. Elsevier.

Dean, N. and A. E. Raftery (2005). Normal uniform mixture differential gene expression detection for cDNA microarrays. *BMC Bioinformatics 6*(1), 173–187.

Dey, D. K., L. Kuo, and S. K. Sahu (1995). A Bayesian predictive approach to determining the number of components in a mixture distribution. *Statistics and Computing 5*, 297–305.

Do, K. A., P. Muller, and F. Tang (2005). A Bayesian mixture model for differential gene expression. *Journal of the Royal Statistical Society, Series C 54*(3), 627–644.

Efron, B., R. Tibshirani, J. D. Storey, and V. Tusher (2001). Empirical Bayes analysis of a microarray experiment. *Journal of the American Statistical Association 96*, 1151–1160.

Freije, W. A., F. E. Castro-Vargas, Z. Fang, S. Horvath, T. Cloughesy, L. M. Liau, P. S. Mischel, and D. F. Nelson (2004). Gene expression profiling of gliomas strongly predicts survival. *Cancer Research 64*, 6503–6510.

Fruhwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis 15*, 183–202.

Gautier, L., L. Cope, B. M. Bolstad, and R. A. Irizarry (2004). affy–analysis of Affymetrix GeneChip data at the probe level. *Bioinformatics 20*(3), 307–315.

Gelfand, A. E., S. Banerjee, and D. Gamerman (2005). Spatial process modelling for univariate and multivariate dynamic spatial data. *EnvironMetrics 16*, 465–479.

Gentleman, R., V. Carey, D. Bates, B. Bolstad, M. Dettling, S. Dudoit, B. Ellis, L. Gautier, Y. Ge, J. Gentry, K. Hornik, T. Hothorn, W. Huber, S. Iacus, R. Irizarry, F. Leisch, C. Li, M. Maechler, A. Rossini, G. Sawitzki, C. Smith, G. Smyth, L. Tierney, J. Yang, and J. Zhang (2004). Bioconductor: open software development for computational biology and bioinformatics. *Genome Biology 5*(10), 1–16. R80.

Hartigan, J. A. (1990). Partition models. *Communications in Statistics - Theory and Methods 19*(8), 2745–2756.

Irizarry, R. A., B. M. Bolstad, F. Collin, L. M. Cope, B. Hobbs, and T. P. Speed (2003). Summaries of Affymetrix GeneChip probe level data. *Nucleic Acids Research 31*(4), e15.

Irizarry, R. A., B. Hobbs, F. Collin, Y. D. Beazer-Barclay, K. J. Antonellis, U. Scherf, and T. P. Speed (2003). Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics 4*, 249–264.

Jong, K., E. Marchiori, G. Meijer, A. V. Vaart, and B. Ylstra (2004). Breakpoint identification and smoothing of array comparative genomic hybridization data. *Bioinformatics 20*, 3636–3637.

Kent, W. J. (2002). BLAT – the BLAST-like alignment tool. *Genome Research 12*(4), 656–664.

Lewin, A., N. Bochkina, and S. Richardson (2007). Fully Bayesian mixture model for differential gene expression: simulations and model checks. *Statistical Applications in Genetics and Molecular Biology 6*, article 36.

Li, M., Y. Wen, Q. Lu, and W. J. Fu (2013). An imputation approach for oligonucleotide microarrays. *PLoS One 8*(3), e58677.

Liu, J. S. (1994). The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem. *Journal of the American Statistical Association 89*, 958–966.

Lucas, J. E., H. N. Kung, and J. T. A. Chi (2010). Latent factor analysis to discover pathway associated putative segmental aneuploidies in human cancers. *PLoS Computational Biology 6*, e1000920.

Marks, J. R., A. M. Davidoff, B. J. Kerns, P. A. Humphrey, J. C. Pence, R. K. Dodge, D. L. Clarke-Pearson, J. D. Iglehart, R. C. Bast, and A. Berchuck (1991). Over-expression and mutation of p53 in epithelial ovarian cancer. *Cancer Research 51*, 2979–2984.

Mayrink, V. D. and J. E. Lucas (2013). Sparse latent factor models with interactions: analysis of gene expression data. *The Annals of Applied Statistics 7*(2), 799–822.

Mayrink, V. D. and J. E. Lucas (2015). Bayesian factor models for the detection of coherent patterns in gene expression data. *Brazilian Journal of Probability and Statistics 29*(1), 1–33.

Miller, L. D., J. Smeds, J. George, V. B. Vega, L. Vergara, A. Ploner, Y. Pawitan, P. Hall, S. Klaar, E. T. Liu, and J. Bergh (2005). An expression signature for p53 status in human breast cancer predicts mutation status, transcriptional effects, and patient survival. *PNAS - Proceedings of the National Academy of Sciences of the United States of America 102*, 13550–13555.

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika 37*(1), 17–23.

Nascimento, F. F., D. Gamerman, and H. F. Lopes (2012). A semiparametric Bayesian approach to extreme value estimation. *Statistics and Computing 22*(2), 661–675.

Newton, M., C. Kendziorski, C. Richmond, F. Blattner, and K. Tsui (2001). On differential variability of expression ratios: improving statistical inference about gene expression changes from microarray data. *Journal of Computational Biology 8*, 37–52.

Page, G. L. and F. A. Quintana (2015, April). Spatial product patition models. arXiv:1504.04489v1.

Parmigiani, G., E. S. Garrett, R. Anbazhagan, and E. Gabrielson (2002). A statistical framework for expression-based molecular classification in cancer. *Journal of the Royal Statistical Society, Series B 64*, 1–20.

Picard, F., S. Robin, M. Lavielle, C. Vaisse, and J. J. Daudin (2005). A statistical approach for CGH microarray data analysis. *BMC Bioinformatics 6*(1), 1–27. 27.

Pollack, J. R., T. Sorlie, C. M. Perou, C. A. Rees, S. S. Jeffrey, P. E. Lonning, R. Tibshirani, D. Botstein, A. L. B. Dale, and P. O. Brown (2002). Microarray analysis reveals a major direct role of DNA copy number alteration in the transcriptional program of human breast tumors. *Proc. of the National Academy of Sciences of the United States of America 99*, 12963–12968.

R Core Team (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.

Roberts, G. O., A. Gelman, and W. R. Gilks (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability 7*(1), 110–120.

Shephard, N. (1994). Partial non-gaussian state space. *Biometrika 81*, 115–131.

Sotiriou, C., P. Wirapati, S. Loi, A. Harris, S. Fox, J. Smeds, H. Nordgren, P. Farmer, V. Praz, B. H. Kains, C. Desmedt, D. Larsimont, F. Cardoso, H. Peterse, D. Nuyten, M. Buyse, M. J. V. D. Vijver, J. Bergh, M. Piccart, and M. Delorenzi (2006). Gene expression profiling in breast cancer: understanding the molecular basis of histologic grade to improve prognosis. *Journal of the National Cancer Institute 98*, 262–272.

Wang, Y., J. G. M. Klijn, Y. Zhang, A. M. Sieuwerts, M. P. Look, F. Yang, D. Talantov, M. Timmermans, M. E. M. V. Gelder, J. Yu, T. Jatkoe, E. M. J. J. Berns, D. Atkins, and J. A. Foekens (2005). Gene expression profiles to predict distant metastasis of lymph-node-negative primary breast cancer. *Lancet 365*, 671–679.

Warren, P., D. Taylor, P. G. V. Martini, J. Jackson, and J. Bienkowska (2007). PANP - a new method of gene detection on oligonucleotide expression arrays. *Proceedings of the 7th IEEE International Conference on Bioinformatics and Bioengineering*.

Wiper, M., D. R. Insua, and F. Ruggeri (2001). Mixtures of gamma distributions with applications. *Journal of Computational and Graphical Statistics 10*(3), 440–454.

Xu, Y., J. Lee, Y. Yuan, R. Mitra, S. Liang, P. Muller, and Y. Ji (2013). Nonparametric Bayesian bi-clustering for next generation sequencing count data. *Bayesian Analysis 8*(4), 759–780.

Yi, Y., J. Mirosevich, Y. Shyr, R. Matusik, and A. L. G. Jr (2005). Coupled analysis of gene expression and chromosomal location. *Genomics 85*, 401–412.